

JOHN N. ...
IN-60-2R
256766
398

HIGH VOLUME DATA STORAGE ARCHITECTURE ANALYSIS

FINAL REPORT

SwRI Project No. 05-3269

Prepared for:

University of Houston Clear Lake
2700 Bay Area Boulevard
Houston, Texas 77058-1096

Subcontract No. 054
RICIS Research Activity No. SE.29

NASA Cooperative Agreement NCC9-16

(NASA-CR-186239) HIGH VOLUME DATA STORAGE
ARCHITECTURE ANALYSIS Final Report
(Southwest Research Inst.) 39 p CSCL 09B

N90-14784

Unclas
63/60 0256766

January 19, 1990



SOUTHWEST RESEARCH INSTITUTE
SAN ANTONIO HOUSTON DETROIT
WASHINGTON, DC DALLAS/FT. WORTH

SOUTHWEST RESEARCH INSTITUTE
Post Office Drawer 28510, 6220 Culebra Road
San Antonio, Texas 78228-0510

HIGH VOLUME DATA STORAGE ARCHITECTURE ANALYSIS

FINAL REPORT

SwRI Project No. 05-3269

Prepared by:
James M. Malik

Prepared for:
University of Houston Clear Lake
2700 Bay Area Boulevard
Houston, Texas 77058-1096

Subcontract No. 054
RICIS Research Activity No. SE.29

NASA Cooperative Agreement NCC9-16

January 19, 1990

Approved:

A handwritten signature in cursive script, reading "Melvin A. Schrader". The signature is written in dark ink and is positioned above a horizontal line.

Melvin A. Schrader, Director
Data Systems Department

TABLE OF CONTENTS

1.0	INTRODUCTION	1
2.0	RESEARCH PERFORMED	1
2.1	List of Contacts	1
2.1.1	Significant Contributors	2
2.1.2	Additional References	4
2.2	Literature Reviews	4
3.0	OPERATIONAL CHARACTERISTICS	8
3.1	SSCC Operational Characteristics	9
3.2	SSCC Operational Requirements	9
3.3	System Characteristics	10
4.0	HIGH VOLUME DATA STORAGE SYSTEMS	11
4.1	Data Systems	11
4.1.1	National Geophysical Data Center	11
4.1.2	IRIS	11
4.1.3	Seismology Data System	11
4.1.4	National Weather Service	12
4.1.5	University of Wisconsin	12
4.1.6	National Center for Atmospheric Research	12
4.1.7	U. S. Geological Survey	12
4.1.8	Harvard	13
4.1.9	Aquidneck	13
4.1.10	Shell Oil Company	13
4.1.11	National Security Agency	13
4.2	Mass Storage Systems	14
4.2.1	Common File System	14
4.2.2	MESA Archival Data Library System	14
4.2.3	Data Facility Hierarchical Storage Manager	14
4.3	Recommendations for In-Depth Analysis	15
4.3.1	University of Wisconsin	15
4.3.2	MESA Archival Data Library System	15
4.3.3	Los Alamos Common File System	15
4.4	Auxiliary Sites	15
4.4.1	Shell Oil Company	15
4.4.2	Aquidneck	16
4.4.3	IBM's Data Storage Products	16

TABLE OF CONTENTS (Continued)

5.0	IN-DEPTH SITE ANALYSIS	16
5.1	University of Wisconsin	16
5.1.1	Site Characterization	16
5.1.2	System Architecture	17
5.1.3	Data Rates/Throughput	20
5.1.4	Archive Capacities	20
5.1.5	Storage Hierarchy and Migration Philosophy	20
5.1.6	Retrieval Capabilities	20
5.2	Mesa Archival Systems, Inc.	21
5.2.1	Site Characterization	21
5.2.2	System Architecture	21
5.2.3	Data Rates/Throughput	22
5.2.4	Archive Capacities	22
5.2.5	Storage Hierarchy and Migration Philosophy	23
5.2.6	Retrieval Capabilities	23
5.3	Common File System	23
5.3.1	Site Characterization	23
5.3.2	System Architecture	24
5.3.3	Data Rates/Throughput	24
5.3.4	Archive Capacities	24
5.3.5	Storage Hierarchy and Migration Philosophy	24
5.3.6	Retrieval Capabilities	25
6.0	EMERGING TECHNOLOGIES	25
6.1	Applicable Technologies	26
7.0	APPLICATION TO SSCC	26
7.1	Archive Configuration	27
7.2	Portability	27
7.3	Proposed Architecture	27
7.4	Shelf Life	28
7.5	Data Access	28
8.0	FURTHER RESEARCH	28
8.1	Core Data Characterization	28
8.2	Data Retrieval	29
8.3	Requirements Analysis	29
8.4	Network Throughput	29
8.5	Application of Database Technology	30
8.6	Design For Long Life	30
8.7	Mass Storage Software	30
9.0	REFERENCES	30

LIST OF FIGURES

FIGURE 1. SSEC McIDAS Configuration	18
FIGURE 2. Archive Recorder Hardware Configuration	19
FIGURE 3. Archive Player Hardware Configuration	19

1.0 INTRODUCTION

This final report documents the effort and findings of Southwest Research Institute (SwRI) in the performance of a High Volume Data Storage Architecture Analysis. This analysis was performed for the National Aeronautics and Space Administration (NASA), Johnson Space Center (JSC) under the NASA Cooperative Agreement NCC9-16, Subcontract No. 054. The results of this analysis will be applied to problems of high volume data requirements such as those anticipated for the Space Station Control Center (SSCC).

2.0 RESEARCH PERFORMED

Prior to the start of the project, SwRI met with Carol Evans, National Aeronautics and Space Administration (NASA) Johnson Space Center (JSC), and Dr. Glen Houston, Research Institute for Computer and Information Systems (RICIS) University of Houston Clear Lake (UHCL), to scope the effort and define project direction.

In the first phase of the project, SwRI performed literature searches and telephone interviews to identify technologies for storing and retrieving large volumes of data.

These initial interviews were instrumental in the identification of potential commercial and Government sites for analysis. SwRI conducted a preliminary analysis of identified sites to select three sites for in-depth analysis. This preliminary analysis was performed based on information gathered during telephone interviews and the review of literature provided by SwRI's contacts or identified via project-specific literature searches.

In an effort parallel to the telephone interviews and literature reviews, SwRI reviewed the operational characteristics anticipated for the Space Station Control Center. This review was limited to information provided by NASA JSC. The Space Station Control Center Level A Requirements were provided by NASA. SwRI also received copies of overheads prepared by MITRE, also under contract with NASA JSC to perform a related study.

In the final phase of the analysis, SwRI visited the three sites selected for in-depth analysis.

The following paragraphs identify SwRI's contacts, documents reviewed by SwRI, and other literature reviewed in the course of this analysis.

2.1 List of Contacts

SwRI made numerous contacts by phone and conducted interviews to support the technology investigation and to identify potential sites for analysis. The following paragraphs contain a list of individuals identified by SwRI

to support its investigation. Individuals are grouped by the office with which each is associated.

2.1.1 Significant Contributors

The individuals listed in this section provided information relevant to SwRI's technology investigation. In combination a wealth of information was provided and has contributed to the success of this investigation.

Satellite Data Processing and Distribution, National Oceanic and Atmospheric Administration (NOAA), Department of Commerce

Bill Callicut	301-763-4640
Dr. Chris Hayden	608-264-5325
Helen Wood	301-763-1564
Bud Booth	301-763-4781
Jack Copan	301-763-1564

University of Wisconsin

Eric Suomi	608-263-6751
William L. Hibbard	608-263-4427

National Aeronautics and Space Administration (NASA)

Strat Laios	301-286-3211
Ron Buch	301-286-9791

National Center for Atmospheric Research (NCAR)

Bernard T. O'Lear	303-497-1268
-------------------	--------------

National Climatic Data Center (NCDC), NOAA, Department of Commerce

Levine Lauritson	301-763-8402
Charles Carpenter	301-763-1372
Rex Snodgrass	704-259-0750
Gus Schembera	704-259-0474
Captain Dropt	301-763-1195

National Geophysical Data Center, NOAA, Department of Commerce

Nettie Bunch	303-497-6150
Ted Habermann	303-497-6472

National Weather Service, NOAA, Department of Commerce

Deanye Lawrence	301-427-7262
Robert Saffold	301-427-7772
Debbie Van de Mark	301-427-7624

National Security Agency

John Davis	301-859-4801
Mark Goldberg	301-859-6555

IRIS

Fumiko Tagima	512-471-0461
Tim Ahern	512-471-0461
Becky Wofford	512-471-0403

Dr. Richard Sailor	617-942-2000
Scott Halbert	505-844-4637
Information Systems, Department of the Treasury	
William Patriara	202-436-6860
Connie Craig	202-436-6565
Pat O'Connor	202-436-6662
Shell Oil Company	
Pat Savage	713-663-2384
U. S. Geological Survey	
Ray Buland	303-236-1506
National Science Foundation	
Marie Zemonkova	202-357-9570
Dr. Michael Foster	202-357-7936
Systems Technology, Information Systems Management, Securities and Exchange Commission	
Eric Malmstrom	202-272-7182
Harvard University	
John Woodhouse	617-495-2637
Greg Williams	203-263-0697
Storage Technology Corporation	
Al Buckland	303-673-3313
Epoch Systems	
Dave Koury	214-387-5277
Exabyte Corporation	
Paule Terrety	713-953-9074
Steve Small	303-442-4333
Mass Store Incorporated	
Gary Smith	301-577-8833
Cray Systems	
David Blaskovich	612-681-3676
Tom Lanzatella	612-681-3354
Paul Rutherford	612-681-3223
Mesa Archival	
Robert I. Smith Jr.	508-842-5336
John W. McIntosh	303-447-1499
Terrence R. D. Rollo	303-447-1499
Mho Salim	303-447-1499

Aquidneck Systems International, Inc.
August David

401-295-2691

2.1.2 Additional References

Due to limited project scope, contact with the following individuals has not been accomplished. SwRI provides this list as a resource to future researchers.

National Center for Super Computer Applications (NCSA) University of Illinois

Barbara Mihalas

Information Processing Division, NOAA, Department of Commerce
Ben Watkins 301-763-5687

National Aeronautics and Space Administration (NASA)

Tom Taylor	301-286-8892 (5520)
Jim Green	(Goddard)
Gary Martin	713-483-9544
Jim Kibler	804-864-5386
Dr. King	301-286-5909

National Climatic Data Center (NCDC), NOAA, Department of Commerce

Bill Burkhart	301-763-4300
Henry Phillips	301-763-5687
Herschel Suits	704-259-0680

Environmental Satellite Data and Information Service, NOAA, Department of Commerce

Irving Perlroth, Data Base Management Division
Bruce Parker, Information Services Division

Internal Revenue Service

John Devlin	202-343-0611
Daniel Capozzoli	202-566-4007
Bill Stalcup	202-343-0611

Trademark and Patent Office

University of Miami Florida

Otis Brown

John Berger 619-534-2889

Dennis Luck 301-688-5065

2.2 Literature Reviews

This section contains bibliographical entries for all literature reviewed by SwRI in support of this effort.

Suomi, Eric W., "The Videocassette GOES Archive System-21 Billion Bits on a Videocassette," IEEE Transactions on Geoscience and Remote Sensing, Vol. GE-20, No. 1, January 1982.

Luck, Dennis R., "The Development of a Modular 10^{14} - 10^{16} Bit Mass Storage Library," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 3.

Manns, Basil, Wilder, Dean, "Interfacing a Mainframe Database Retrieval System With An Optical Disk Image Storage System," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 7.

Nelson, Marc, Kitts, Davis L., Merrill, John H., Harano, Gene, "The NCAR Mass Storage System," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 12.

Kempster, Linda S., Martin, John B., "In Search of: NASA Space Data Storage Solutions," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 27.

Halford, Robert J., "Mass Storage Mechanization for Cray Computer Systems," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 52.

Bedoll, R. F., "FMS - File Management System at Boeing Computer Services," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 66.

Burgess, John, "Virtual Library System: A General Purpose Mass Storage Archive," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 72.

DeVries, John, "NFS - An Approach to Distributed File Systems in Heterogeneous Networks," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 77.

Burke, James J., Hu, Paul Y., "The Optical Data Storage Center," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 89.

Itao, Kiyoshi, Yamaji, Akihiko, Hara, Shigeji, Izawa, Nobuyoshi, "Magneto-Optical Mass Storage System with 130mm Write-Once Disk Compatibility," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 92.

Bessette, Oliver, "High Performance Optical Disk for Mass Storage Systems," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 98.

Funkenbusch, A. W., Rinehart, T. A., Siitari, D. W., Hwang, Y. S., Gardner, R. N., "Magneto-optics Technology for Mass Storage Systems," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 101.

Kurtz, Clark, "Development of a High-Capacity, High-Performance Optical Storage System," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 107.

Larson, David, D., Young, James R., Studebaker, Thomas J., Kraybill, Cynthia L., "StorageTek 4400 Automated Cartridge System," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 112.

Mitsuya, Y., Takanami, S., Koshimoto, Y., Sato, I., "8.8-GBYTE Capacity Magnetic Storage System," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 118.

Muraco, Paul F., "D-1 Magnetic Tape Mass Storage Application," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 124.

Oelschlaeger, Jon R., "Mass Storage Systems: An Applications View," Digest of Papers Eighth IEEE Symposium on Mass Storage Systems, May 1987, p. 135.

Weiss, James R., Riegler, Guenter R., "Managing Data in the Great Observatory Era," Information Systems Newsletter, Pasadena, California, Issue 16, April 1989, p. 22.

Domchick, Hal, Naughton, Patricia, "NSESOC Converts Library to New 3480 Tape Cartridge System," Information Systems Newsletter, Pasadena, California, Issue 16, April 1989, p. 37.

Green, James L., "What Can We Learn From an Online Archive?," NSSDC (National Space Science Data Center) News, Vol. 4, Nos. 3/4, Fall/Winter 1988, n.p., p. 2.

Krishnaswamy, Sumant, King, Joseph H., Kayser, Susan, "International Sun-Earth Explorer Data Will Be Archived Over Three-Year Period," NSSDC (National Space Science Data Center) News, Vol. 4, Nos. 3/4, Fall/Winter 1988, n.p., p. 11.

McClanahan, Scott, "Magneto-Optical Disks," MOSL Newsletter, NASA/Johnson Space Center, Houston, Texas, Volume 1 Number 1 (September 1989), p. 9.

O'Lear, Bernard T., Kitts, David L., "Optical Mass Data Storage II," A Reprint From The Proceedings of SPIE-The International Society for Optical Engineering, San Diego, California, 18-22 August 1986

Miller, Stephen W., "Mass Storage System Reference Model: Version 2.0," Menlo Park, California, May 1987.

O'Lear, Bernard T., Choy, Joseph H., "Software Considerations in Mass Storage Systems," Reprinted from Computer Magazine, Los Alamitos, California, July, 1982, p. 36.

O'Lear, Bernard T., Choy, Joseph H., "Optical Device Interfacing for a Mass Storage System," Reprinted from Computer Magazine, Los Alamitos, California, July, 1985, p. 24.

Hartman, Berl, "OLTP On The VAXCLUSTER," DEC Professional, January 1988.

Hibbard, William, Santek, David, "Visualizing Large Data Sets in the Earth Sciences," Computer Magazine, Los Alamitos, California, August, 1989, p. 53-57.

Collins, Bill, Devaney, Marjorie, Kitts, David, "Profiles in Mass Storage: A Tale of Two Systems," Computing and Communications Division, Los Alamos, New Mexico, and National Center for Atmospheric Research, Boulder, Colorado, n.d.

2.3 Document Reviews

This section contains a list of documentation for hardware and software systems reviewed by SwRI.

"SSP 30261 Architectural Control Document - Data Management System, Rev. B, 02/19/88, NASA Space Station Program Office," Space Station Control Center (SSCC) Level A Requirements Original Issue, NASA/Johnson Space Center, October 1989, p. 2-1.

"Optical Archiving System Product Description, System O A S 150," Aquidneck Systems International, Inc., N. Kingstown, RI, n.d.

"Testing Space Shuttle Main Engines," Concurrent Computer Corporation, Tinton Falls, NJ, n.d.

"Data Acquisition and Analysis," Concurrent Computer Corporation, Tinton Falls, NJ, n.d.

"Telemetry," Concurrent Computer Corporation, Tinton Falls, NJ, n.d.

Henize, John, "Understanding Real-Time UNIX," Concurrent Computer Corporation, Houston, Texas.

Atlas, Alan, Blundon, Bill, "Time To Reach For It All," Reprinted with permission from UNIX REVIEW, n. p., January 1989.

"UNISYS With Concurrent Awarded NEXRAD," Concurrent Computer Corporation, Customer Focus, Houston, Texas, Spring 1988.

"HPD368F, Fixed Disk System," Concurrent Computer Corporation, Customer Focus, Oceanport, NJ, n.d.

"OS/32 & MTM, Real-Time Operating System and Multi-Terminal Monitor," Concurrent Computer Corporation, Oceanport, NJ, n.d.

"3280E MPS, Multiprocessor System," Concurrent Computer Corporation, Tinton Falls, NJ, n.d.

"3212 Computer System," Concurrent Computer Corporation, Oceanport, NJ, n.d.

"Epoch-1 InfiniteStorage Server," Epoch Systems, Inc., Marlborough, MA, 1988.

"New Data Storage Strategies For High Performance Workstations," Epoch Systems, Inc., Marlborough, MA, 1988.

"Epoch Systems Announces World's Highest Capacity Workstation Server," Epoch Systems, Inc., Marlborough, MA, Oct., 1989.

"EXB-8200 8mm Cartridge Tape Subsystem, Interface User Manual" Exabyte Corporation, Boulder, Colorado, November 1988.

"EXB-8200 8mm Cartridge Tape Subsystem, Product Specification" Exabyte Corporation, Boulder, Colorado, November 1988.

"EXB-8200 8mm Cartridge Tape Subsystem, Product Overview" Exabyte Corporation, Boulder, Colorado, July 1987.

"IBM Data Facility Hierarchical Storage Manager, Version 2 Release 5.0, General Information," IBM Corporation, Tuscon, Arizona, Sixth Edition, July 1989.

"Space Science and Engineering Center," Space Science and Engineering Center, Madison, Wisconsin, Revised September, 1987.

"Application Brief, University of Wisconsin-Madison, IBM Academic Information Systems," International Business Machines Corporation, Milford, CT, November 1986.

"The Data Library," Mesa Archival Systems, Inc., Boulder, Colorado, Release 1.2, July 1989.

"The NCAR Mass Storage System," NCAR Scientific Computing Division, University Corporation for Atmospheric Research, 1988.

3.0 OPERATIONAL CHARACTERISTICS

The following paragraphs describe operational characteristics of high volume data storage and retrieval systems.

3.1 SSCC Operational Characteristics

On October 19, 1989, SwRI met with representatives from NASA JSC and RICIS UHCL in a pre-project meeting. In this meeting, NASA provided the following characteristics for the SSCC.

- Anticipated continuous data rate of 50 gigabytes per day
- Time stamped data
- Multi-user environment
- Distributed system
- High-volume, long-term archive
- Priority given to data integrity and minimizing data loss
- Retrieval times under 5 minutes for near-real time data
- Relaxed retrieval times for older data
 - . 3 day old data should be retrievable from on-line or near-line storage (e.g. automated tape library or optical disk jukebox).
 - . 3 month old data may be stored on off-line media which should be available on-site to allow the data to be loaded onto an on-line media within 24 hours.
 - . 3 year old data may be stored on off-line media which resides off site.

3.2 SSCC Operational Requirements

SwRI reviewed the Space Station Control Center Level A Documentation for operational characteristics. This document provided few specifics relative to this analysis. High level requirements stated in this document which have bearing on this and subsequent efforts are summarized below.

Core data processing and archiving is one of seven areas of responsibility defined for the SSCC.

The SSCC will consist of data acquisition and transmission, data distribution, data processing, data storage and retrieval, and support system elements.

The Level A document defines eight other ground elements which interface with the SSCC. SSCC interfaces are not limited to these ground elements. Requirements governing external interface support are also defined.

The SSCC will be housed in a five-story building with approximately 106,000 square feet. Other characteristics of the facility are provided. However, allocations to each of the areas of responsibility are not specified.

Performance shall be measured against valid requirements in terms of the time required by SSCC to accept, process, and return correct output for a user input. Performance should be sufficient to guarantee mandatory, highly desirable, and routine functions without risk to crew or to success of the mission.

Growth capability goals are specified to reflect the need for the capability to incorporate changes in existing or future technology and to address the needs to increase capacity or functionality.

Similar requirements are outlined for commonality, reusability, interoperability, flexibility, automation, tailorability, and human factors.

The SSCC must be designed to provide security, privacy, integrity protection, disclosure protection, and access control. The SSCC shall restrict commanding operations to designated locations according to command sensitivity level, and user and location authorization.

Reliability, maintainability and availability goals are described. Of particular interest to those tasks of recording and archiving core data is the requirement which states that the data capture function shall be maintained in the event of system failure. The data capture function shall have a maximum allowable outage of one minute over a one week period.

Specific requirements relative to the development of data storage and retrieval functions are summarized below.

The SSCC shall provide thirty minute access to both flight and ground data which is one year old or less.

The SSCC shall permanently archive selected flight and ground data which is greater than one year old and retrieve this data within 24 hours of request.

3.3 System Characteristics

In the course of its analysis, SwRI identified operational characteristics common among the mass storage systems reviewed. These characteristics are summarized below.

The data archive typically operates as a single node in a heterogeneous computing network.

The data archive must support data retrieval requests from numerous computing platforms in the network.

Data retrieval requests are typically bursty in nature as opposed to the continuous nature anticipated for acquisition and storage.

It is difficult to characterize data usage and therefore the nature of the data requests. Hence, the storage format must be generic to support flexible retrieval services.

4.0 HIGH VOLUME DATA STORAGE SYSTEMS

In an effort to identify potential commercial and government sites for analysis, SwRI made numerous contacts by phone. Initially SwRI called individuals associated with data processing, information systems, or system development offices in several government agencies. The nature of the investigation was described and each individual was asked to describe systems in their domain which were used for high volume data storage. SwRI also contacted several hardware vendors. Most of the individuals contacted provided names of individuals performing related research or having responsibility for data storage systems. In some cases system documentation was solicited by SwRI.

4.1 Data Systems

In the course of the investigation of technologies for storing and retrieving large volumes of digital data, SwRI identified commercial and government data systems for analysis. These systems are described in the following paragraphs.

4.1.1 National Geophysical Data Center

Nettie Bunch with the Information Services Division provided information about their data storage system. Data from satellites and earthquake stations are received in various formats. This data is reformatted and written to off-line media including magnetic tapes and Write Once Read Many (WORMS) optical disks. Individual's data managers maintain the archive index.

4.1.2 IRIS

The IRIS data center in Austin, Texas utilizes an IBM mainframe to archive seismology data. This data is gathered on a system in Albuquerque and transmitted to Austin for archival and subsequent distribution to universities. The data is transmitted on tapes which are loaded onto the IBM using a SUN microcomputer and a hyperchannel link. On-line capacity is 8 gigabytes. Data is retrieved by day, time, and geographical location. Distribution tapes are generated by the SUN system. Data retrieval is a slow process which may require hours to complete. Application of WORM technology and a jukebox library is being investigated.

4.1.3 Seismology Data System

A VAX cluster running VMS is used to gather, process and archive seismology data collected at numerous earthquake stations. The data is

received in various formats, primarily on magnetic tape. It is staged onto magnetic disks, processed and then written to WORM disks using an Aquidneck controller for the Sony jukebox. The Sony jukebox provides 150 gigabytes of near-line storage. Distribution volumes are created as data from each time period arrives from the stations. The distribution volumes are created on magnetic tape. IRIS receives its data via these distribution volumes. Backup and some distribution is performed using Exabyte's 8 mm helical cartridge tape system.

4.1.4 National Weather Service

Robert Saffold of the National Weather Service described the development of NEXRAD. NEXRAD is a system which will employ approximately 150 remote sites to collect data and store it on WORM disks. These disks will be sent to a central location in Ashville operated by the National Climatic Data Center (NCDC). He indicated that the WORM disks are simply stored on racks. He also indicated that data retrieval has not been defined.

4.1.5 University of Wisconsin

The University of Wisconsin is under contract to archive satellite data for the National Oceanic and Atmospheric Administration (NOAA). A Sony video system has been adapted for data archival. Data is stored on a Sony pneumatic 3/4" video tape which has a capacity of approximately 10 gigabytes. The data is received from each satellite at a rate of 1.7 megabits/second for 18 minutes every half hour. This translates into approximately 11 gigabytes per day. Data collection and archival has been accomplished for up to three satellites. Retrieval is supported by search information which has been recorded on one of the tape's audio tracks. This information includes the satellite identifier, scan number, and Julian day. [Suomi]

4.1.6 National Center for Atmospheric Research

Bernard O'Lear at the National Center for Atmospheric Research (NCAR) provided information about their mass storage system which has a capacity of eleven terabytes. This system uses IBM drives, custom software, and Storage Technology products. As a result of our phone conversation, Mr. O'Lear has provided documentation for the NCAR Mass Storage System and numerous publications regarding High Volume Data Storage.

4.1.7 U. S. Geological Survey

Ray Buland with the U.S. Geological Survey in Colorado is in the process of acquiring a system similar to that used in Albuquerque. Data is acquired at a rate of 60 megabytes per day. This rate is expected to increase five fold over the next three years as new stations are brought on-line. As in Albuquerque, the data is received on cassette tapes, staged into magnetic disks and archived on WORMS. This data is used to create a final volume. The data is retrievable by day, time, and station. Requests are typically of two types; long time periods for one station or

short time periods for multiple stations. Mr. Buland indicated that acquisition of a system to retrieve real-time data is planned.

4.1.8 Harvard

John Woodhouse has set up a data archival and management system at Harvard for seismology data. This system employs a Sony jukebox with a capacity of 165 gigabytes. The jukebox is controlled by the Aquidneck controller. Presently the archive is being migrated from a Data General platform to a Sun platform. In the interim the jukebox is mechanically switchable between the Data General and the Sun. Data is received on tape, processed, and stored on the jukebox. Programs have been developed to read and extract segments of data.

4.1.9 Aquidneck

August David of Aquidneck has offered to host site visits at several sites employing the Aquidneck controller in combination with WORM disks in a jukebox. Mr. David offered visits to the Houston Chronicle which has a two jukebox system and Woodlands Geophysical which aids geologist and geophysicist with a range of interpretation and archiving needs. He also provided contacts at NASA JSC, and Texaco.

4.1.10 Shell Oil Company

Pat Savage of Shell Oil Company manages a system which employs 3480 technology to archive seismic data collected in the field. Mr. Savage indicated that he has 2 million reels of data in archive. He also expressed a high degree of confidence that the mass storage requirements for core data from space station could be met with proven 3480 technology. He stated that this technology is very reliable and offers high performance and wide acceptance.

4.1.11 National Security Agency

In the course of its investigation, SwRI was directed to the National Security Agency (NSA) by several individuals. SwRI contacted two individuals at NSA. However, both were reluctant to provide specific information about NSA systems. SwRI was told that NSA was developing a system using IBM 3480 technology interfaced to a VAX environment. NSA is involved in efforts to force the development of a mass storage device with a capacity of 1,000 terrabits, transfer rates of 100 megabits/second, and useable directories.

Mr. Goldberg provided helpful insight which should influence the design of any high volume data storage system. First, he indicated that use of optical disk technology should be limited to systems which must provide its user with control over random access retrieval of the archived data. Second, he pointed out that requirements for media with a 30 year life assumes the hardware used to playback the data has a 30 year life. He

emphasized this by pointing out that it would be quite difficult to acquire or maintain hardware capable of reading any media used to store data in the 1960's.

While SwRI cannot recommend any site for further analysis, we do recommend that NASA establish contacts at NSA. At minimum this should allow free information exchange. Ideally, NASA could cooperate with NSA to bring forth technologies to satisfy common requirements.

4.2 Mass Storage Systems

In the performance of this investigation, SwRI has identified mass storage systems (MSS) which are either commercially available or have been installed at several sites. These systems are described in the following paragraphs.

4.2.1 Common File System

The Los Alamos Common File System (CFS) is a file storage and file management system that serves heterogeneous computing networks. It provides a centralized file storage and file access capability for all machines in the Los Alamos Integrated Computing Network (ICN). The CFS provides in excess of seven terabytes of storage for machines in the ICN. The CFS software has been installed in at least seventeen other computing sites. The ICN consists of supercomputers, general purpose computers, scientific workstations, and personal computers. The CFS provides archival storage, storage for inactive files, and backup services. [Collins]

4.2.2 MESA Archival Data Library System

MESA Archival's Data Library System (DLS) is a complete file archive management system designed for high performances and ease of use in a networked computing environment. The DLS is an implementation of the Institute for Electrical and Electronics Engineers (IEEE) Computer Society Reference Model of Mass Storage. The DLS may be attached to most commercial computers. Its network access server provides the interface to commercially available network software. This system supports a hierarchy of storage devices. It locates the most active files on the fastest access devices and the least active files on lower cost-per-bit devices.

4.2.3 Data Facility Hierarchical Storage Manager

IBM offers a line of products which provide system-managed storage. These products work together to determine data placement, automatically manage data availability, performance and space, and relieve users of data management details. IBM's Storage Management Products provide an integrated approach toward an IBM system-managed storage environment.

4.3 Recommendations for In-Depth Analysis

The following paragraphs document SwRI's selection of sites for in-depth analysis.

4.3.1 University of Wisconsin

SwRI recommends selection of the Geostationary Operational Environment Satellite (GOES) videocassette archive system for in-depth analysis. This system incorporates real-time data acquisition, high volume storage, and a unique concept for maintaining index information to facilitate retrieval. The storage media is long-life, high density, and low cost. Since the system has been on-line since the early 1980's and incorporates custom leading edge technology, SwRI believes useful insight into the life cycle of a state-of-the-art system would be gained.

4.3.2 MESA Archival Data Library System

Because the NCAR Mass Storage System (MSS) follows the IEEE Computer Society Reference Model for Mass Storage Systems, it is an excellent candidate for in-depth analysis and a site visit. However, Mr. O'Lear has suggested that any visit would have to be in mid-January or later. He also requested early notice of any planned visits. However, SwRI understands that MESA Archival's Data Library System (DLS) also follows the IEEE model and evolved from NCAR's MSS. SwRI recommends in-depth analysis of the MESA Archival DLS.

4.3.3 Los Alamos Common File System

The Los Alamos Common File System is the third system recommended by SwRI for in-depth analysis. Despite the fact that this system does not perform data acquisition, SwRI believes that it is a good candidate for in-depth analysis. The Los Alamos CFS provides in excess of seven terabytes of data for a heterogenous computing network. It also supports file movement with burst rates of 50 Megabits/second. [Collins]

4.4 Auxiliary Sites

SwRI had planned to augment the information assimilated during in-depth analysis of the selected sites with visits to auxiliary sites; however, project scope and schedule prevented SwRI from visiting these sites. The following paragraphs describe the auxiliary sites.

4.4.1 Shell Oil Company

The mass storage system at Shell Oil Company is another good candidate for in-depth analysis. However, only three sites were to be selected for in-depth analysis. Because Shell Oil Company is located in Houston and is readily accessible, a site visit to further analyze the technologies employed for application to the Space Station Control Center environment is recommended.

4.4.2 Aquidneck

August David of Aquidneck offered site visits to several sites utilizing the Aquidneck controller for optical disk storage. SwRI cannot recommend any of the Aquidneck sites on their own merit. However, a visit to one of the Houston installations would provide an opportunity to review a system employing optical disk technology.

4.4.3 IBM's Data Storage Products

SwRI did not recommend in-depth analysis of IBM's data storage products. However, SwRI believes there is merit in reviewing a commercially available product.

5.0 IN-DEPTH SITE ANALYSIS

The following paragraphs document the results of the in-depth analysis for the three sites selected.

5.1 University of Wisconsin

On January 8, 1990, SwRI met with Eric Suomi at the Space Science and Engineering Center (SSEC), University of Wisconsin. Mr. Suomi described the use of an adapted video recorder to record high-speed digital data from the Geostationary Operational Environmental Satellite (GOES) series of satellites. He also provided demonstrations of the Man computer Interactive Data Access System (McIDAS).

SwRI originally planned the site visit to review the videocassette archive exclusively. However, after seeing both the GOES videocassette archive and the McIDAS systems, SwRI believes it is appropriate to discuss both systems as they are related to the archival and interactive access of GOES data.

5.1.1 Site Characterization

The Space Science and Engineering Center (SSEC) at the University of Wisconsin is a multidisciplinary research and development center. SSEC's stated mission follows:

- Atmospheric studies of Earth and other planets,
- Interactive computing, data access and image processing, and
- Space flight hardware development and fabrication.

SSEC developed the videocassette archive system to record high-speed digital data from the GOES satellites. [Suomi] SSEC collects geostationary satellite data in digital format on customized videocassette

tapes and has done so since 1978. The digital equivalent of one hundred Libraries of Congress has been collected and archived.

McIDAS is an interactive tool which facilitates the combination of data access and processing power of the computer with reasoning, judgement, and pattern recognition skills of the user.

McIDAS is a powerful data management and analysis tool which supports:

- Meteorological research,
- Operational weather forecasting, and
- Education.

McIDAS features include:

- Real-time data
- Interaction (user-guided computer processing)
- Weather analysis tools
- User adaptable applications
- Potential for growth through new data sources and applications.

McIDAS is a design philosophy as well as a set of hardware and software. McIDAS allows the user to access tremendous amounts of raw data and apply applications to generate information. Because McIDAS is an integrated set of tools, it is constantly evolving. This evolution feeds itself as users develop custom applications by integrating existing McIDAS features to solve new problems. These solutions may in turn evolve into McIDAS tools and become a part of the core system.

Although McIDAS supports the analysis of data from numerous sources, SwRI has limited its review to GOES satellite data.

5.1.2 System Architecture

The hardware platform for the McIDAS systems includes a Model 4381 IBM Mainframe, peripheral storage in excess of 33GB, and tape drives recording at either 6250 or 1600 bits per inch. This computer platform is integrated into a configuration which includes antennas, a network of remote computers, ingestors, and archive playback hardware. Figure 1 depicts the McIDAS architecture.

The videocassette archive system consists of an adapted video recorder which has been integrated with an encoder and power supply. Figure 2 depicts the archive recorder hardware configuration. A similarly adapted unit has been integrated with a controller, video monitor, and decoder to provide playback. Figure 3 depicts the archive player hardware configuration. [Suomi]

SSEC McIDAS CONFIGURATION

(JUNE 1987)




FIGURE 1. SSEC McIDAS Configuration

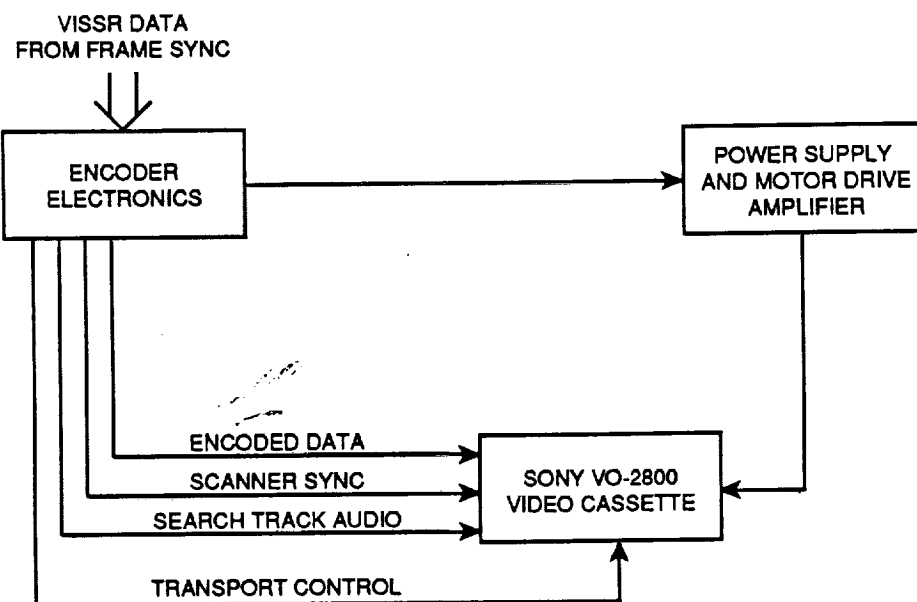


FIGURE 2. ARCHIVE RECORDER HARDWARE CONFIGURATION

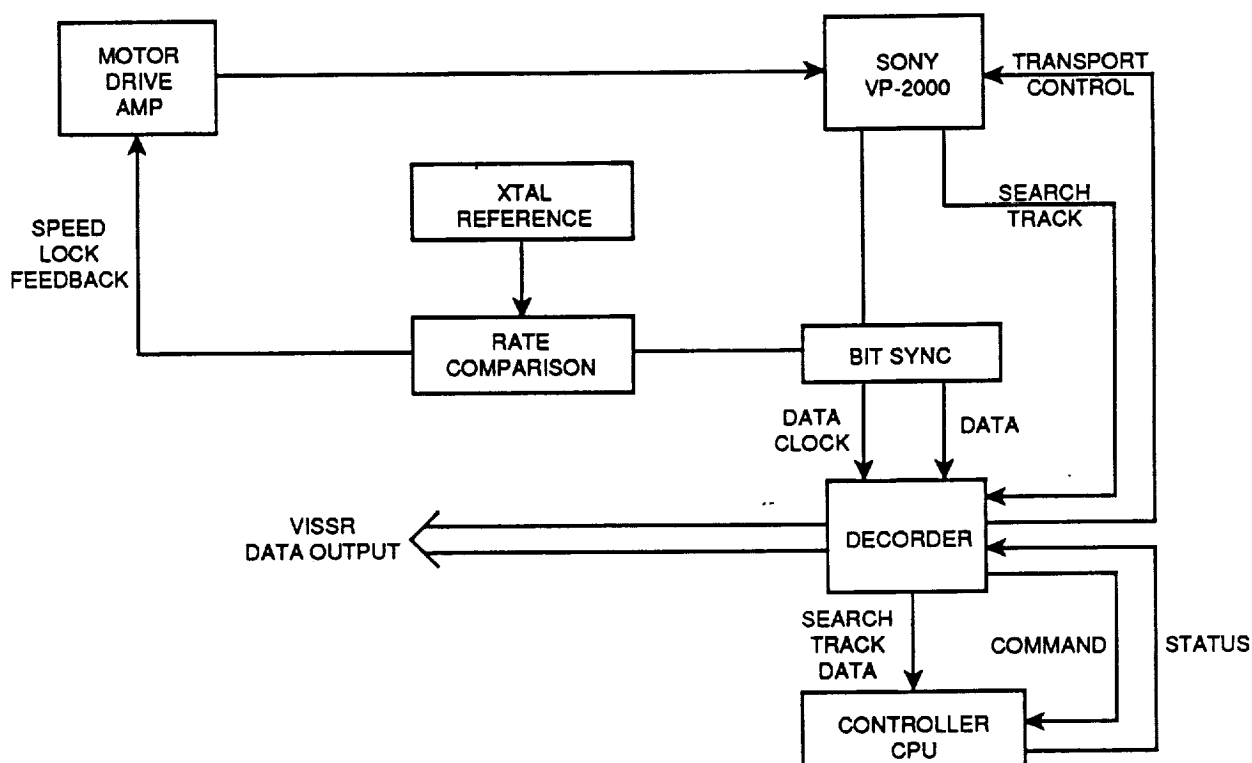


FIGURE 3. ARCHIVE PLAYER HARDWARE CONFIGURATION

5.1.3 Data Rates/Throughput

McIDAS receives in excess of 5 GB/day. Because only the most recent four to six images are saved only 592 MB of this data is maintained on-line. Prior to the loss of the second GOES satellite in early 1989, in excess of 10 GB was received each day. The McIDAS system also receives data from two other satellites and other ground based equipment making the total daily retrieval rate approximately 15 GB.

Presently, the GOES videocassette archive receives and archives approximately 19 GB each day. The satellite transmits 2.1136 Mbits/second for 25 minutes of every half hour. The system has archived as much as 33 GB/day received from 3 satellites transmitting 1.7472 Mbits/second each for 18 minutes every half hour. [Suomi]

5.1.4 Archive Capacities

The McIDAS system is configured with 33 GB of direct access storage. Approximately 15 GB is used for temporary storage of satellite data and data from ground based equipment. The McIDAS system does not archive data.

The videocassette archive has been on-line since 1978. The data archive contains an estimated 40 terabytes of GOES satellite data.

5.1.5 Storage Hierarchy and Migration Philosophy

The McIDAS system does not archive data. It maintains the most recent four to six images on direct access storage. Older images are purged from the system.

The videocassette archive is not hierarchical. Data is recorded directly onto the archive media. It is anticipated that the data will remain on the video media throughout its useful life.

5.1.6 Retrieval Capabilities

The McIDAS system is an open system which allows the user to develop custom applications for retrieving and analyzing data. The core system provides analysis tools which access the data files maintained on the direct access storage devices. The data is stored in a generic file structure designed to allow easy data access from utility programs and to eliminate redundant sorting/editing routines.

The videocassette tapes must be mounted in player hardware to perform data retrieval. Typically the playback system is operated manually. However, the playback system does provide some automated search capabilities.

The videocassette archive player hardware can be used to upload data into the McIDAS system.

5.2 Mesa Archival Systems, Inc.

On January 11, 1990, SwRI met with John McIntosh, Terrence Rollo, and Mho Salim of Mesa Archival Systems in Boulder, Colorado. These individuals described the Data Library System (DLS) which has been commercialized and marketed by Mesa. The DLS is a commercialized version of the NCAR Mass Storage System. Mesa's DLS is an implementation of the Institute for Electrical And Electronics Engineers (IEEE) Computer Society Reference Model of Mass Storage.

5.2.1 Site Characterization

Mesa's DLS product is in its infancy with existing installations numbering less than five. Planned installations number between three and five. For the purpose of this discussion, the installation at NCAR will be referenced. NCAR provides computer power and data storage needed by atmospheric researchers for extensive modeling and data analysis.

5.2.2 System Architecture

The DLS is a software product which consists of three major software components:

The Data Library Control Program (DLCP) is the core software of the DLS. It runs as an application under the IBM operating system MVS/XA. The DLCP processes user requests to store and retrieve files and to manage directories. It automatically performs system administration tasks such as media management and validation of data integrity. The DLCP utilizes a Master File Directory which maintains directory information for all files in the archive.

The Network Access Server operates at the presentation and application layers of the ISO model to provide an interface to commercially available network software which operates at the session and transport layers. This server software also runs under the IBM operating system MVS/XA.

The Data Library Access software runs on each user computer to allow users to store and retrieve files with standard commands from a wide variety of computers and operating systems on the network.

The Data Library Processor is the computer which hosts the Data Library Control Program and the Network Access Server software. Data archives and the Master File Directory are maintained on Data Library Processor storage peripherals. The Data Library Processor can be connected to the Data Library Access software on user computers through a variety of commercial data networks. The Data Library System does not manage any files on storage devices attached to user computers.

5.2.3 Data Rates/Throughput

An estimated 96 gigabytes of data is transferred between the NCAR MSS and user computers each day.

Mr. John McIntosh of Mesa Archival provided the following network performance estimates for Mesa's DLS.

The transfer of a 10 GB file from the user computer to the Data Library Processor (DLP) using a single Ultraset path to an HPPI channel on an IBM 3090 would require 6 minutes based on an average sustained transfer rate of 30 MB per second.

The transfer of a 10 GB file from the user computer to the DLP using a single HYPERchannel path to a block multiplexor channel on an IBM 3090 is about 135 minutes based on an average sustained transfer rate of 1.25 MB per second.

The transfer of a 10 GB file from the user computer to the DLP using a single Ethernet path to a block multiplexor channel on an IBM 3090 is about 8,400 minutes (six days) based on an average sustained transfer rate of 20 KB per second.

Mr. McIntosh emphasized that network data transfer performance is dependent on many factors including the network configuration, the number of network paths available, the volume of network traffic, the command to data ratio, data block sizes, the specific CPU configuration, the operating system, the network adapter, the network protocol, and the workload characteristics of the systems involved. He also indicated that no definitive studies are available that address network performance in any controlled environment.

5.2.4 Archive Capacities

Mr. O'Lear estimates that NCAR's archival system provides access to 9 Terabytes of data stored on 58,000 IBM 3480 tape cartridges. Each cartridge has a capacity of 200 MB. Cartridge utilization is estimated to be 81%.

Mr. McIntosh has proposed an architecture to provide archive capabilities for core data to be received at an estimated rate of 50 gigabytes per day. The proposed architecture is illustrated in Attachment A and summarized below:

- IBM 3090 110J
- 120 GB IBM 3380 disk
- 6 IBM 3480 cartridge tape transports
- Cartridge tape robotic system

5.2.5 Storage Hierarchy and Migration Philosophy

The NCAR Mass Storage System disk farm is constantly monitored to determine the best methods to tune the system to increase the disk "hit rate". The disk hit rate has been increased from 38% to 66% by adding partitions for smaller bitfiles and automatically staging bitfiles which have been read twice in a five-day period to disk.

When a file is received by Mesa's DLS it is temporarily placed on one of the DLS disks. If the user does not access the file within a customer-specified time period or if disk space must be freed, the system automatically migrates the file to the archival devices. The DLS uses disk to buffer file transfers to and from archival devices (e.g. 3480 cartridges). The file movement process is transparent to the user.

5.2.6 Retrieval Capabilities

Mesa's DLS is a file archive system. User files are stored in the archive as bitfiles. Files are retrieved from the archive and transferred to the user's computer in response to requests initiated at the user computer.

5.3 Common File System

On December 28, 1989, SwRI met with Paul Rutherford of Cray Research Mendota Heights, Minnesota. Mr. Rutherford described their use of the Common File System (CFS). Mr. Rutherford also described Cray Products which can be integrated to provide a high performance mass storage system.

CFS is a commercial product marketed by General Atomics. It has been integrated into numerous computer networks world-wide. In "Profiles in Mass Storage: A Tale of Two Systems", Collins, Devaney, and Kitts describe the Los Alamos Common File System and the NCAR Mass Storage System. SwRI has supplemented the information gained from its site visit with information from this article.

5.3.1 Site Characterization

The computing network at Cray Research, supports the development of Cray products as well as other scientific research. Significant compute power is provided by the network of supercomputers, general purpose computers, and workstations. Individual workstations provide users with a platform for research and development. The supercomputers and general purpose computers provide computing horsepower and storage for researchers. Researchers tend to store their most relevant data on the workstation while allowing the less frequently used data to remain on other network storage devices. The network provides a permanent store which is a data storage resource available to researchers. In this environment, CFS is used to archive data migrated from the permanent store.

The Los Alamos Integrated Computing Network is a scientific computing network of many different machines running eight different operating systems. File storage, output processing, data import/export, access

control, job control and other services are provided by network support servers. Network supercomputers are used interactively for program development, job setup, execution of short jobs, and output analysis. At night, production jobs are run in batch. CFS is used to store job, input, and output files for the production jobs. The CFS provides centralized file storage and file access for network servers and machines. [Collins]

5.3.2 System Architecture

In the configuration reviewed at the Cray site, the CFS software resides on an IBM 3090 with 40 Gigabytes (GB) of on-line disk storage and multiple tape drives. The tape drives are not supported by an automated loading system. The IBM 3090 is connected to a HYPERchannel high speed network. Multiple Cray systems are connected to the same high speed network. Approximately 500 SUN Workstations access the high speed network via an ETHERNET local area network. Approximately 2500 Sun Workstations worldwide access the high speed network via a wide area network.

CFS is integrated into the Los Alamos Integrated Computing Network to provide centralized file storage and access. Collins describes the network as a large scientific computing network of supercomputers, general purpose computers, scientific workstations, and personal computers. CFS utilizes the Los Alamos File Transport System and gateways to receive and transmit user request/responses and files. [Collins]

5.3.3 Data Rates/Throughput

Data rates at the Cray site are limited by the I/O bandwidth of the 3090 and are estimated to be one Megabyte per second. In the month of November, in excess of 40,000 file transfers were processed with total I/O in excess of 125 GB. Fifty-four percent of these requests were satisfied from disk and forty-six percent from tape.

Collins reports that data transfer exceeds 50 GB per day in the Los Alamos installation. [Collins]

5.3.4 Archive Capacities

At the Cray site, total CFS system storage is approximately 327 GB. Of this, 14 GB resides on disk with a total capacity of 40 GB and 313 GB resides on tape.

Collins reports total storage in excess of 7 Terabytes (TB) at the Los Alamos installation with a growth rate of over two TB per year. [Collins]

5.3.5 Storage Hierarchy and Migration Philosophy

At the Cray site, a data migration facility front ends the CFS archive which is totally hidden from the user in the current configuration. Presently, data migrates to permanent store on the network via NFS. The data management facility (DMF) moves data files from the permanent store to the CFS archive. Small files are written to disk and large files are

written to tape. The DMF manages data retrieval from the permanent store. While user requests initiate retrieval from the permanent store, DMF invokes CFS to retrieve data files migrated to the CFS archive.

In the Los Alamos installation the user is given more control and flexibility but is required to be more knowledgeable. The user must take explicit action to store, retrieve, delete, convert and backup files. [Collins]

5.3.6 Retrieval Capabilities

At the Cray site, the data management facility retrieves data from the CFS archive to satisfy file transfer requests for files removed from the permanent store.

In the Los Alamos installation, retrieval is initiated by the user.

6.0 EMERGING TECHNOLOGIES

SwRI has identified numerous hardware components which can be applied to the problem associated with high volume data storage. Applicable computer platforms include super minicomputers, high-end mainframes, and low-end supercomputers. SwRI believes the critical factors in selection of a computer platform, from most significant to least significant, are:

- I/O bandwidth
- Addressable memory
- CPU performance

Numerous storage platforms were also identified in the course of SwRI's analysis. Applicable platforms are listed and characterized below:

- IBM 3480 compatible tape cartridges: 3480 technology is reliable, offers high performance, and has wide acceptance. It is supported by numerous software and hardware products.
- High-performance, high-capacity tape systems like the EXABYTE EXB-8200 CTS offer low cost storage. This product utilizes advanced helical scan technology to provide high recording densities and storage capacities.
- Optical disk technology offers high density storage and random access of data. While its cost per bit and access times are less favorable than other high density technologies, it may find applicability because of its random access capability. Anticipated improvements should dramatically lower cost per bit while improving access times.

- Magnetic disks will be used to facilitate near real-time access to data and to buffer archive media I/O.

Numerous network options exist. An array of hardware platforms, protocols, standards, and products exist and must be evaluated. SwRI anticipates continued improvements in network communications technology.

6.1 Applicable Technologies

NASA should monitor developments for the following emerging technologies which may be applicable in the high volume data storage environment for SSCC.

- Optical tape technology has not evolved as expected. Yet, it promises high-density storage at a lower cost per bit than optical disk. If this technology evolves to meet current expectations, it should provide another media option for the archive system.
- D-2 is an emerging tape format standard. At this time, no D-2 products exist, and few are under development. However, in the future, D-2 products should offer high density and low cost per bit storage.

7.0 APPLICATION TO SSCC

During the preliminary analysis phase of the project, when SwRI was performing telephone interviews and literature reviews, options and products seemed almost limitless. Many of our contacts had high expectations for technologies like optical disk or tape. However, as SwRI started to identify systems with high volume data storage components, we discovered the dominant media is magnetic tape.

SwRI understands that the systems reviewed either as sites selected for in-depth analysis or via phone interviews with system administrators, developers, or users, are mature systems. Hence, the dependence on "mundane" technologies is understandable. However, the focus of this investigation was on technologies in use today in systems with high volume storage requirements.

SwRI believes that new and evolving technologies will impact the development of high volume data storage systems. Further, SwRI believes NASA should cultivate the development of high-density, low-cost media and anticipate use of new technologies to meet the high volume data storage requirements of SSCC. However, SwRI cautions NASA to avoid trendy products. NASA should pursue products which are both widely accepted and supported and are based on accepted standards.

The following paragraphs discuss application of the analysis results to the high volume data storage requirements of SSCC.

7.1 Archive Configuration

SwRI believes there is merit in dedicating a machine for the purpose of archival. This machine should be configured with a hierarchy of storage devices. Archive software which provides the functionality of the following IEEE-CS MSS modules should reside on this machine:

- Bitfile Server
- Storage Server
- Bitfile Mover
- Name Server
- Site Manager

7.2 Portability

Collins reports that CFS progressed across hardware platforms with minimal change due to extensive use of MVS software. [Collins] In the design phase for the archive software, consideration should be given to adaptability to permit use of new storage media.

7.3 Proposed Architecture

SwRI solicited proposed architectures from contacts at sites selected for in-depth analysis. John McIntosh, Mesa Archival, provided a model for a hypothetical data archiving system. This proposed architecture is included as Attachment A.

Paul Rutherford, Cray Research, provided the foundation for a very high speed file server with a four terabyte capacity. This proposed architecture is included as Attachment B.

Attachment C is a proposed architecture developed by SwRI. This proposed architecture is at a high level by design. SwRI does not believe the defined requirements or the scope of this project yield themselves to a more detailed proposal. This architecture is designed to reflect the following characteristics:

The applicability of a medium speed (6 megabits/second) recorder should be evaluated. This recorder would be used to record the core data before it is processed. It could be used to validate the storage processor, back up the archive platform in the event of failure, or provide the media for long-term archival.

The archive platform, whether central or distributed, should provide both storage processing and retrieval processing. The retrieval processor should transmit only the data required by the user. This will minimize network traffic which will result in improved response times.

Current to three-day old data should reside on a direct access media.

The archive storage should be hierarchical. Lower-cost, slower-access devices should be buffered with faster-access storage media.

7.4 Shelf Life

SwRI recognizes that different shelf life requirements exist for the archive media and the archived data. SwRI recommends that shelf life requirements for the media and the data be expressed separately. It has been pointed out that a media with a shelf life of thirty years is useless if the hardware required to read data from the media is obsolete after ten years.

7.5 Data Access

The McIDAS system is an open-ended system designed to expand to meet the needs of its users. Data is stored in generic file structures designed to allow easy data access from utility programs and to eliminate redundant sorting/editing routines. The core system provides data management and analysis tools. In this system the data is stored in a format which facilitates retrieval by researchers/users.

Webster defines an archive as a place in which public records or historical data is preserved. The design of the SSCC high volume data storage system should provide for flexible data access rather than efficient archival. The concept of an open-ended system which provides tools to facilitate data access should be evaluated for applicability.

8.0 FURTHER RESEARCH

The following paragraphs identify topics for continued research to support the acquisition of a high volume data storage system for SSCC.

8.1 Core Data Characterization

In the course of the analysis, the questions, "What does the data look like?" and "What is the format of the data being archived?" were asked by SwRI and their contacts. SwRI understands that the data originates from Space Station Freedom and is limited to digital data.

SwRI has speculated that the data can be processed to generate fixed format records with well defined field content. If this is the case, time dependent relational tables could be used to store the data in a manner which would facilitate retrieval.

Even though the 50 GB/day retrieval rate is expected to be continuous, SwRI has assumed that numerous time intervals will govern the frequency at which individual values are transmitted.

These questions must be answered before meaningful descriptions of storage formats can be defined.

8.2 Data Retrieval

Characterization of the data retrieval requests must also be achieved before meaningful descriptions of storage formats can be defined. SwRI understands that retrieval capabilities better suited to the needs of the users than the playback mode used in other NASA systems is desired. However, retrieval requirements, the nature of the data requests, and retrieval data rates are undefined.

SwRI believes that filtering of the data must be performed on a platform with high speed access to the data archive. This will help to minimize network traffic which will in turn improve response time.

8.3 Requirements Analysis

It may seem that characterization of the data and retrieval requests has little to do with a technology investigation. However, without this information, the formulation of conceptual designs is meaningless. This must be tempered with the understanding that SwRI believes computing platforms, network configurations, and archive media exist today and will evolve to meet the requirements of this high volume data storage application. However, SwRI does not believe that software to support this volume of data is readily available.

Supported by the discussions in the previous paragraphs, SwRI recommends that NASA work to develop descriptions of the data to be stored and operational concepts for the data retrieval subsystem.

8.4 Network Throughput

John McIntosh, President of Mesa Archival, and SwRI recognize that data transfer performance in a network environment is dependent on many factors:

- Network configuration,
- Number of network paths available,
- Volume of the network traffic,
- Command to data ratio,
- Data block sizes,
- CPU configuration,
- Operating system,
- Network adapter,
- Network protocol, and
- System workload characteristics.

Availability of studies which address network performance in any controlled environment is minimal. Research which will support the definition of the hardware and software platform to meet the communication requirements of the archive system is crucial to successful development.

8.5 Application of Database Technology

Considering the anticipated 30-year life of the Space Station Freedom, it is unlikely that all data storage and retrieval requirements can be anticipated prior to the development of the high volume data storage system. Given this, adaptability should be a high-priority system design goal. Relational database systems have proven to be quite adaptable.

A second database technology which may prove useful in the design of the data storage system is object oriented databases. The applicability of relational or objected oriented database systems should be researched.

8.6 Design For Long Life

NASA has a history of designing for long life. Recent trends including software portability and the application of standards during software development should be extended. Additional research concerned with achieving long life for software systems should be pursued.

8.7 Mass Storage Software

SwRI failed to identify any software platform which provided functionality similar to the anticipated requirements for the SSCC high volume data storage system. SwRI did identify a few data systems capable of storing and retrieving files from large capacity archive systems. However, these systems concerned themselves with user-identified units of data (files), not with the data content of those files. Even at this level, the directories maintained by this system are adequate at best. Continued research in the following areas is critical to the successful development of a high volume storage system for SSCC:

- Directory structures which support the anticipated data volumes.
- Directory structures which support hierarchical archive configurations for hierarchies with more than two storage media.
- Data dictionaries to support user-defined access of data from the high volume data storage system.

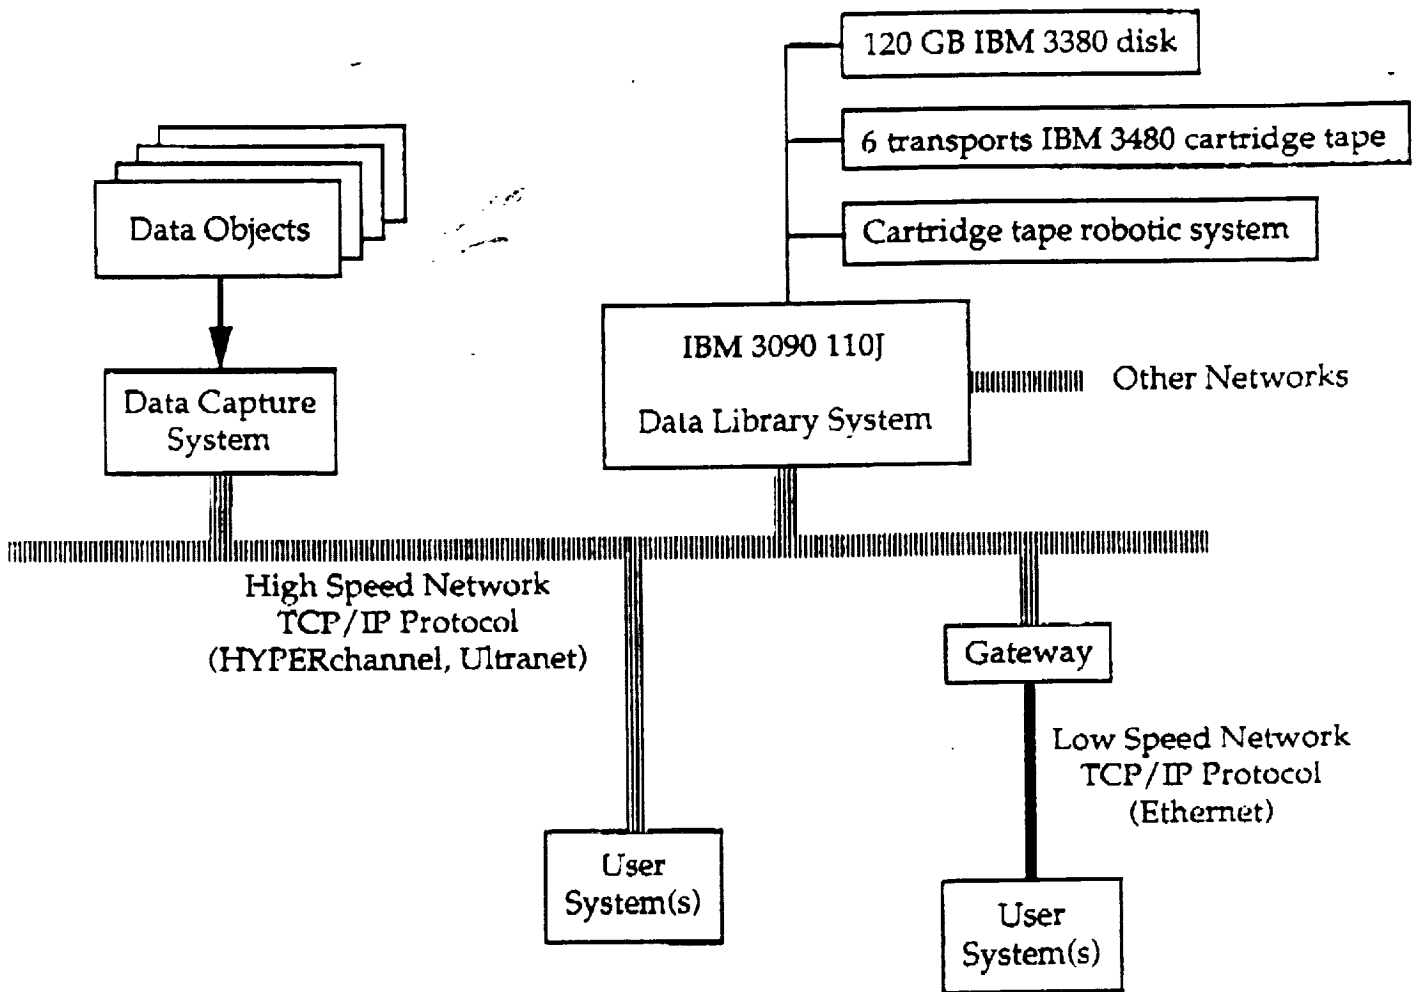

9.0 REFERENCES

Suomi, Eric W., "The Videocassette GOES Archive System-21 Billion Bits on a Videocassette," IEEE Transactions on Geoscience and Remote Sensing, Vol. GE-20, No. 1, January 1982.

Collins, Bill, Devaney, Marjorie, Kitts, David, "Profiles in Mass Storage: A Tale of Two Systems," Computing and Communications Division, Los Alamos, New Mexico, and National Center for Atmospheric Research, Boulder, Colorado, n.d.


ATTACHMENTS

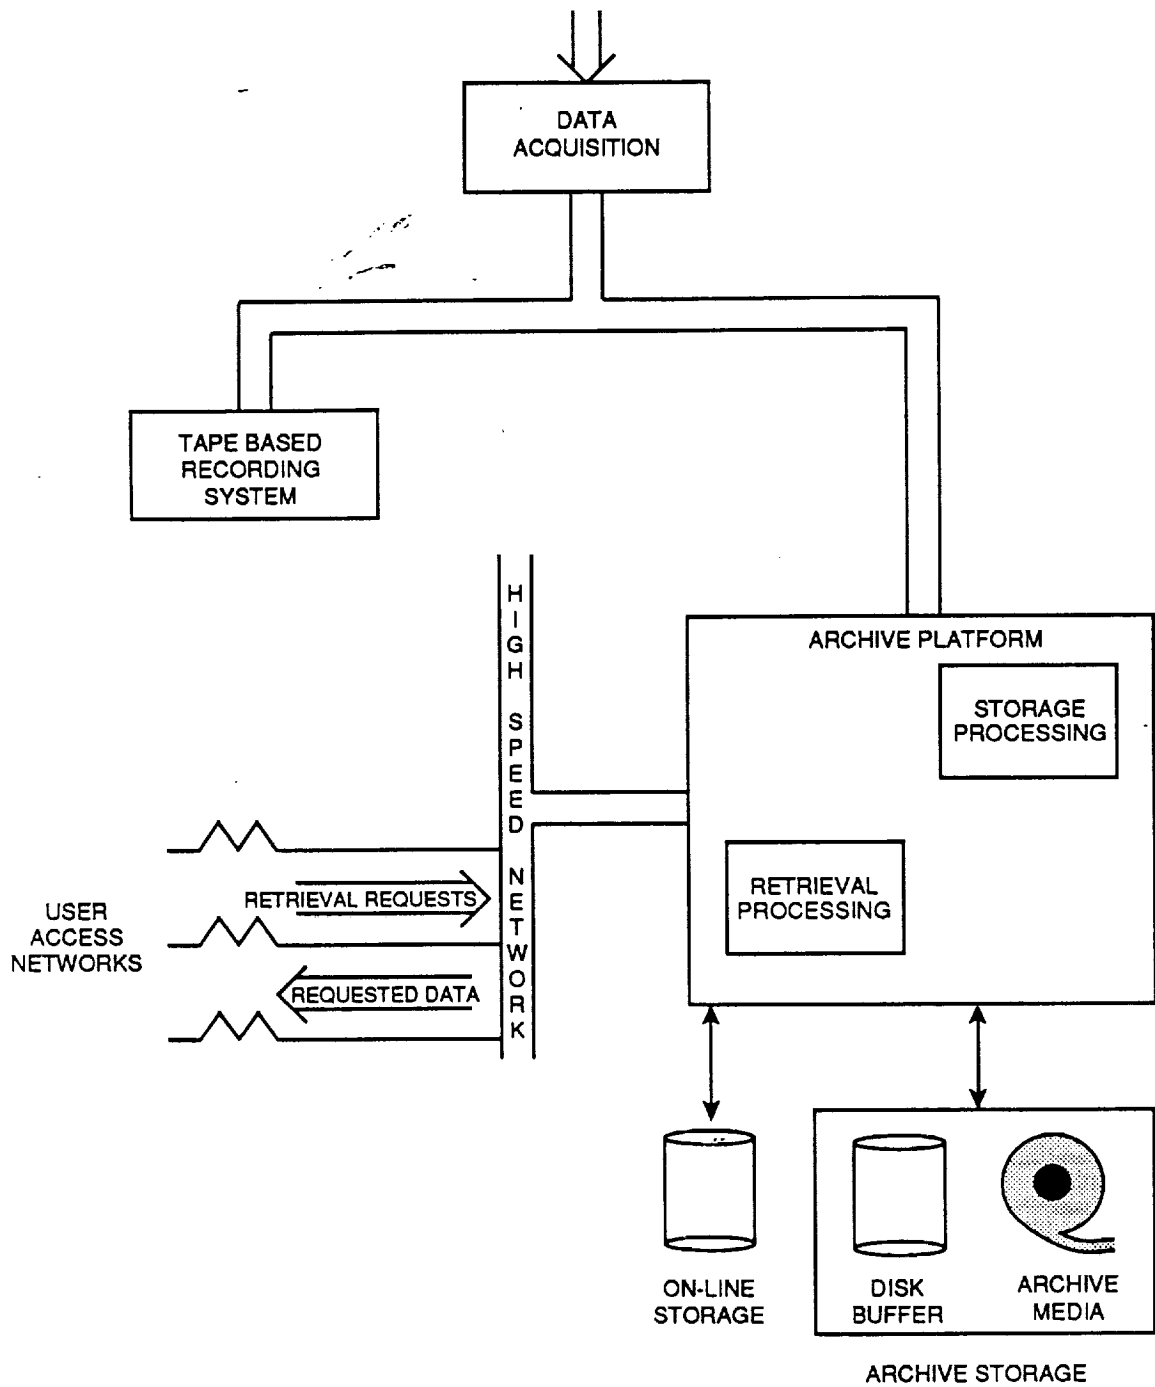
Mesa Archival



Hypothetical Data Archiving System

Storage Superserver Hardware





ATTACHMENT C
PROPOSED ARCHITECTURE FOR HIGH VOLUME DATA STORAGE

